gsiso.ai
The Agentic Intelligence Fabric

PHASE 0 / KICKOFF PACKET

Phase 0 Kickoff

60-day plan: founding hires, dev infra, design-partner LOIs.
May 9 — July 7, 2026.

OWNER	STATUS	PHASE
Gaurav Sisodia	In progress	0 of 5

Contents

Goal by July 7, 2026: 4 founding engineers signed, 3 design-partner LOIs in hand, dev infra running, Sprint 1 spec frozen. Phase 1 begins July 1 with Sprint 1 (MCP server skeleton + ROS 2 adapter).

Wedge thesis. Physical AI Bridge plus EU AI Act compliance certification. Position as the Red Hat / Palo Alto Networks for agents — governance, identity, audit, and policy as the durable moat.

SECTION 01

Phase 0 Tracker

Week-by-week plan from May 9 to July 7, 2026. Updated every Friday by the owner. Status flips to red when a week's must-ship items slip past Friday EOD.

Phase 0 — 60-Day Kickoff Tracker

Window: May 9 → July 7, 2026 (M0–M2 of the delivery plan) Owner: Gaurav Sisodia Goal at end of Phase 0: 5-person team hired, dev infra running, 3 design-partner LOIs signed, Phase 1 spec frozen.

North-star metrics for Phase 0

Metric	Target by Jul 7	Status
Founding engineers signed (offers accepted)	4 of 4	0/4
Design-partner LOIs signed	3 of 3	0/3
Dev infra: 4 repos + CI green + Terraform dev env up	Done	repos [x], CI/TF pending
Phase 1 spec frozen	Document signed off	Sprint 1 spec drafted
Seed round: term sheet in hand	Yes	not started

Week-by-week plan

Week 1 — May 9–15 (this week)

Theme: Setup + first outreach wave

- [x] Vercel deploy live at gsiso.ai + gsiso.com
- [x] GitHub org bootstrapped — 4 private repos with READMEs and CI scaffolds
- [ ] Set up `gaurav@gsiso.ai` and `hiring@gsiso.ai` email forwarding (Squarespace → Gmail or Fastmail)
- [ ] Post 4 JDs on LinkedIn, Hacker News "Who's Hiring?" thread (June 1), Robotics Career Network, Triplebyte
- [ ] Send first 3 design-partner outreach emails (Recursion, Siemens, Two Sigma)
- [ ] Send 5 follow-up cold emails to robotics PhD network (LinkedIn search)
- [ ] Open AWS account in `gsiso-platform` org; set up SSO; us-east-1 + eu-west-1 enabled
- [ ] Open Cloudflare account; transfer DNS for `gsiso.ai` from Squarespace (optional but recommended for granular DNS control)

Week 2 — May 16–22

Theme: Infra MVP + first interviews

- [ ] Terraform dev environment in AWS us-east-1: VPC, EKS cluster (1 node group), RDS Postgres, Elasticache Redis
- [ ] OPA policy engine deployed in dev EKS as a sidecar pattern reference
- [ ] HashiCorp Vault dev cluster running (1 node, dev mode acceptable for now)
- [ ] First-round interviews with Hire #1 and Hire #2 candidates
- [ ] LOI template lawyer-reviewed and ready to send to design partners

Week 3 — May 23–29

Theme: Spec freeze + offer wave 1

- [ ] Phase 1 Sprint 1 spec frozen (see `specs/sprint-01-mcp-skeleton.md`)
- [ ] First 2 founding-engineer offers extended (Distributed Systems + ROS 2)
- [ ] Recursion + Siemens replies → second meeting scheduled
- [ ] Vault production setup in eu-west-1 + us-east-1 with HSM-backed root keys
- [ ] First test agent DID minted end-to-end (manual, not scheduled)

Week 4 — May 30–Jun 5

Theme: Lab procurement + LOI #1

- [ ] First LOI signed (target: Recursion or Insitro)
- [ ] OT-2 lab arm ordered ($12K · 4–6 week lead time)
- [ ] UR10e cobot ordered ($80K · 6–8 week lead time)
- [ ] NVIDIA DGX workstation ordered for Isaac Sim ($200K · 4 week lead time)
- [ ] Lab space LOI signed (Bay Area: South SF, Burlingame, or Hayward — robot-safe flooring + 3-phase power)
- [ ] Hires #1 and #2 sign offers; start dates negotiated for late June / early July
- [ ] D&O insurance quoted; product liability coverage scoped

Week 5 — Jun 6–12

Theme: Hire #3 + Hire #4 + LOI #2

- [ ] Second LOI signed (target: Siemens or Boston Dynamics)
- [ ] Hire #3 (Security/Crypto) offer extended
- [ ] Hire #4 (Design Engineer) offer extended
- [ ] Trust Ledger v0 design doc complete (ed25519 + Merkle chain + Postgres schema)
- [ ] Sprint 2 spec drafted (Trust Ledger alpha)
- [ ] Seed pitch deck v1 drafted (use delivery-plan.pdf as primary appendix)

Week 6 — Jun 13–19

Theme: Onboarding prep + LOI #3

- [ ] Third LOI signed (target: Two Sigma or Citadel)

- [ ] Hires #3 and #4 sign offers
- [ ] Onboarding doc set written: dev env setup, repo conventions, security posture, on-call rotation
- [ ] First seed investor meetings scheduled (target list of 12 funds; aim for 6 first meetings)
- [ ] OT-2 arrives — install in lab

Week 7 — Jun 20–26

Theme: Hires arrive (early ones)

- [ ] Hires #1 and #2 start (some may negotiate later — minimum 1 starts this week)
- [ ] Sprint 0 retrospective: team alignment on Phase 1 sprint plan
- [ ] First robot in lab — OT-2 brought up with Opentrons SDK
- [ ] Seed term sheet conversations advancing with 2–3 funds

Week 8 — Jun 27–Jul 3

Theme: Sprint 1 readiness

- [ ] All 4 founding hires onboarded or imminent start dates
- [ ] Sprint 1 kickoff readiness check — every prereq from `specs/sprint-01-mcp-skeleton.md` met
- [ ] Phase 0 retrospective written
- [ ] Term sheet decision on lead investor

Buffer week — Jul 4–7

Theme: Phase 1 launch readiness

- [ ] Phase 1 begins July 1 (Sprint 1 dates: Jul 1–11 per delivery plan §4)
- [ ] First sprint demo scheduled for Jul 11

Daily cadence

- Mon 9:00 PT — outreach session (2 design-partner emails sent, 3 follow-ups)
- Tue/Thu — candidate interviews (block calendar)
- Wed 14:00 PT — investor meeting block (when round opens)
- Fri 15:00 PT — Phase 0 weekly review (15 min, async OK if solo); update this tracker
- Sat AM — deep work on the next-up spec doc

Risks for Phase 0 specifically

Risk	Mitigation
Founding engineer #1 (distributed systems) takes > 8 weeks to find	Cast wide on day 1 — LinkedIn + 3 recruiters + warm network. Pre-screen with a 30-min architecture conversation, not a coding test.
No design partner LOI by Jun 30	Drop fees to zero for first anchor in each vertical; offer co-design rights at GA; expand outreach list from 3 → 8
Lab space falls through	Have 2 backup options under LOI before committing; consider SF vs East Bay tradeoff
Cash runway visibility weakens	Open seed conversations by Week 5 (Jun 8) — lead-time on term sheets is 4–6 weeks

Links

- Delivery plan PDF — full 18-month plan
- Site — public marketing
- GitHub org: github.com/sisodiabhumca — repos: gsiso-platform, gsiso-mcp-bridge, gsiso-console, gsiso-infra
- Vercel: gsisodia-projects/gsiso-ai

SECTION 02

Dev Environment Setup

Repo bootstrap, CI scaffold, secret management, and local-stack quickstart for the four founding engineers. Everything below should be runnable on day one of M0.

Dev Environment Setup — Phase 0 (May 2026)

Owner: Gaurav Sisodia (until Hire #1 starts) Goal: Everything an engineer needs on day one is provisioned and reproducible.

What's already done

- [x] Domain: `gsiso.ai` (canonical) + `gsiso.com` (redirect) on Vercel
- [x] Marketing site + console prototype + docs deployed
- [x] GitHub org: 4 private repos created — `gsiso-platform`, `gsiso-mcp-bridge`, `gsiso-console`, `gsiso-infra`
- [x] Repo READMEs, .gitignore, CI scaffolds, and ADR-001 (tech stack lock) committed

What needs doing this week

1. Email + comms (1 hour)

- [ ] Set up Squarespace email forwarding for `@gsiso.ai`:
  - `gaurav@gsiso.ai` → personal Gmail
  - `hiring@gsiso.ai` → personal Gmail
  - `partners@gsiso.ai` → personal Gmail
  - `legal@gsiso.ai` → personal Gmail
- [ ] (Optional) Buy Google Workspace — $7/user/month — for proper send-as. Worth doing once Hire #1 is signed.
- [ ] Slack workspace: free tier OK at this stage. Channels: `#general`, `#engineering`, `#robotics`, `#partners`, `#hiring`, `#random`.

2. AWS account (2 hours)

- [ ] Create new AWS account under `gsiso-platform` business identity
- [ ] Enable AWS Organizations + 4 sub-accounts: `dev`, `staging`, `prod-us`, `prod-eu`
- [ ] Configure SSO via AWS IAM Identity Center (no IAM users for humans)
- [ ] Enable us-east-1 (primary) + eu-west-1 (AI Act residency)
- [ ] Set up budget alerts: $5K/mo dev, $15K/mo staging, $50K/mo prod (will scale up)
- [ ] Enable CloudTrail + GuardDuty in all regions
- [ ] Buy Reserved Instances at quarterly review — not yet

3. Terraform (1 day)

- [ ] In `gsiso-infra` repo, create `terraform/envs/dev/` with:
  - VPC (10.0.0.0/16, 3 AZs, public + private subnets)

  - EKS cluster (1.30, 1 managed node group, 3 × t3.large)
  - RDS Postgres 16 (db.t4g.medium, single-AZ for dev, pgvector extension enabled)
  - Elasticache Redis 7 (cache.t4g.small)
  - S3 buckets: `gsiso-dev-artifacts`, `gsiso-dev-receipts-cold`
- [ ] Run `terraform apply` — verify everything comes up green
- [ ] Save state to S3 backend with DynamoDB locking
- [ ] Document: estimated dev cost ~$800/mo (acceptable)

4. Vault (4 hours)

- [ ] HashiCorp Vault dev cluster in EKS dev — Helm chart, single replica, dev mode acceptable
- [ ] Test: store and retrieve a dummy secret via CLI
- [ ] Plan: production Vault clusters in eu-west-1 and us-east-1 with HSM (AWS CloudHSM or KMS-backed) — schedule for Week 2 of Phase 1

5. Observability scaffolding (2 hours)

- [ ] Create Sentry account on Team plan ($26/mo) — projects for `platform`, `console`, `mcp-bridge`
- [ ] Create Grafana Cloud Free Tier — connect to OpenTelemetry Collector
- [ ] ClickHouse Cloud trial signup — defer to Sprint 2 when audit volume justifies it

6. Vendor accounts to open

- [ ] Linear (project management) — 1 user free, $8/user/mo as team grows
- [ ] Vanta — for SOC2 — schedule sales call for Week 4
- [ ] Stripe — open account, even if not yet billing
- [ ] Clerk — auth for the future console; free tier sufficient for now
- [ ] Anthropic, OpenAI, Google AI Platform — API keys for development; route through LiteLLM
- [ ] LiteLLM Cloud trial — Sprint 1 dependency

7. Legal + financial structure

- [ ] Confirm Delaware C-Corp formed (or use Stripe Atlas — $500, ~5 days)
- [ ] EIN obtained, Mercury or Brex business banking opened
- [ ] Initial cap table in Carta or Pulley — founder shares, option pool ~12% reserved for first 5 hires
- [ ] Standard YC SAFE (post-money, $20M valuation cap recommended for seed) ready for first investor conversations
- [ ] Outside counsel retained — Cooley, Gunderson, or Wilson Sonsini all acceptable for seed-stage AI/robotics

---

Daily / weekly hygiene

- Push something to a repo every day (any of the 4) — even a doc update. Activity history matters for hires.
- One investor coffee per week starting Week 5 (warm intros first; cold pings only after 3 warm meetings).
- Weekly Friday review: update `00-PHASE-0-TRACKER.md`, commit it, push.

Estimated Phase 0 spend (May–June)

Line item	Cost
AWS dev workloads	~$1,600
Vendor SaaS (Sentry, Grafana, LiteLLM, Slack, Linear)	~$400
Email + DNS (Squarespace + maybe Workspace)	~$50
Founder salary (2 months at $175K)	~$29,000
Outside counsel + Stripe Atlas	~$8,000
Recruiter retainer (1 firm, 3 roles)	~$15,000
Total Phase 0 burn	~$54,000

The bulk of Phase 0 spend is the founder's runway and legal — not infra or hardware. Hardware procurement begins Week 4 when LOI #1 is in hand.

SECTION 03

Sprint 1 Spec

MCP server skeleton plus ROS 2 adapter, with OPA policy decisions under 50ms p99. Two-week sprint: July 1 to July 11, 2026 — the first executable artifact of the platform.

Sprint 1 Spec — MCP Server Skeleton + ROS 2 Adapter Scaffolding

Sprint window: Jul 1–11, 2026 (Phase 1, Sprint 1) Owner: Founding Engineer #2 (Robotics / ROS 2) Reviewer: Gaurav Sisodia Demo date: Friday Jul 11, 15:00 PT — engineer walkthrough (~30 min)

Goal

Stand up the MCP server skeleton and the ROS 2 adapter scaffolding so that the Claude/GPT/Gemini agents can issue a *dummy* robot command, have it evaluated by the OPA policy engine, and receive an allow/deny decision in under 50 ms. No real robot, no Isaac Sim yet — pure plumbing.

By the end of this sprint, the path from `agent intent` → `MCP server` → `ROS 2 action stub` → `OPA policy decision` → `audit log entry` works end-to-end on every push.

In scope

1. MCP server stub — Python 3.12, FastAPI surface, single endpoint: `POST /actions/dispatch`. Accepts `{agent_did, robot_class, action, params}`. Returns `{decision, receipt_id, latency_ms}`.

2. ROS 2 Jazzy in Docker — base image with `ros:jazzy-ros-base`; ROS 2 daemon running; `colcon build` working in the container. The MCP server runs *outside* the ROS container in dev (Phase 1) and talks to it via a local message bus stub.

3. OPA integration — OPA running as a sidecar in the dev EKS cluster (or local Docker for laptop dev). One example Rego policy enforced: deny any command where `params.force_threshold > 50N` (placeholder for safe-stop).

4. Audit log entry — for every dispatch, a JSON row written to Postgres `audit_events` table. Schema must support sprint-2's signed receipts without migration. Columns: `id, ts, agent_did, robot_class, action, params, decision, latency_ms, prev_hash, signature` (signature/prev_hash NULL in S1, populated in S2).

5. CI pipeline green — `gsiso-platform` and `gsiso-mcp-bridge` repos must have GitHub Actions running ruff + pytest on every PR. Branch protection on `main`: 1 approving review + CI green.

6. Dev environment runbook — written in `gsiso-platform/docs/setup.md`. New engineer should be able to clone repos, run `make dev`, and have the full stack (Postgres, Redis, OPA, MCP server) running in < 15 min.

Explicitly out of scope

- Real robot connection (Sprint 4 for OT-2)
- ed25519 signing of receipts (Sprint 2)
- WebRTC telemetry (Sprint 4)
- Isaac Sim integration (Sprint 3)
- Multi-tenant Postgres schema (Phase 2)
- Production-hardened OPA policies (incremental)

Acceptance criteria

#	Criterion	How to verify
1	`POST /actions/dispatch` returns 200 with valid payload	curl + manual test in CI
2	Dummy command with `force_threshold: 60` returns `decision: deny`	pytest
3	Dummy command with `force_threshold: 30` returns `decision: allow`	pytest
4	End-to-end p99 latency < 50 ms on dev hardware (M2 Mac or equivalent)	Locust load test, 100 RPS for 60 s
5	Audit row written to Postgres `audit_events` with all non-signature fields populated	pytest with testcontainers
6	CI pipeline runs ruff + pytest, green on every PR	GitHub Actions
7	Setup runbook reproducible — confirmed by Hire #1 (Distributed Systems) following it solo	Manual handoff test

Architecture sketch (v0)

```
+-----------------+ +-----------------+ +-------------+ | Agent | HTTP | MCP Server | HTTP | OPA | | (Claude/etc) |------->| (FastAPI) |------->| (Rego) | +-----------------+ +-----------------+ +-------------+ | | SQL v +------------------+ | Postgres | | audit_events | +------------------+
```

Future: Step from MCP Server to a ROS 2 action client (Sprint 3+); insert ed25519 signing between OPA decision and audit write (Sprint 2).

Tasks (suggested breakdown)

Task	Owner	Estimate
Repo skeleton + pyproject + CI for gsiso-platform	Hire #1	0.5 day
FastAPI app w/ `/actions/dispatch` + Pydantic models	Hire #1	1 day
Postgres schema + Alembic migration for `audit_events`	Hire #1	0.5 day
OPA local + first Rego policy	Hire #1 + Gaurav	1 day
ROS 2 Jazzy Docker base image	Hire #2	1 day
ROS 2 action client stub in MCP server	Hire #2	1 day
`make dev` Compose file (Postgres + Redis + OPA + MCP)	Hire #1	1 day
pytest suite — happy path + deny path + latency check	Hire #2	1 day
`setup.md` runbook	Hire #2	0.5 day
Demo prep — terminal recording + slide w/ architecture diagram	Gaurav	0.5 day

Demo script (Jul 11)

1. Show: clone `gsiso-platform` and `gsiso-mcp-bridge`, run `make dev`. (3 min)
2. Show: `curl -X POST localhost:8000/actions/dispatch -d @safe.json` → `{"decision": "allow", ...}` (1 min)
3. Show: Same with `unsafe.json` (force 60N) → `{"decision": "deny", ...}` (1 min)
4. Show: Postgres `SELECT * FROM audit_events ORDER BY ts DESC LIMIT 5` — both events present (2 min)
5. Show: Locust load test screenshot — p99 < 50ms (2 min)
6. Show: GitHub Actions green badges on both repos (1 min)
7. Q&A and discussion of Sprint 2 (ed25519 signing) (20 min)

Definition of done

A second engineer (or Gaurav) can clone both repos on a fresh laptop, follow `setup.md`, and reproduce all 7 acceptance criteria within 30 minutes.

If the demo on Jul 11 fails any acceptance criterion, the sprint extends by one week per delivery-plan §4 sprint conventions. No exceptions.

SECTION 04

Founding Hires

Four roles. Two start in M0 (distributed systems, ROS 2 robotics). Two start in M1 (security/crypto, design engineer). Compensation bands are inclusive of equity ranges; letters of intent close within seven days of verbal yes.

Founding Engineer — Distributed Systems

Location: Bay Area, CA — Hybrid (3 days on-site) Employment type: Full-time Reports to: Gaurav Sisodia, Founder Hire sequence: #1 — Month 0 Compensation: $180,000–$220,000 base + 1.0–1.5% equity

Why now — the company in 60 seconds

gsiso.ai is building the operating layer that sits above open protocols like MCP and A2A, and below application-layer SaaS. The product is an agentic intelligence fabric: a distributed runtime that schedules, routes, and governs thousands of concurrent AI agents across robotics, pharma, manufacturing, and capital markets. Two dateable, defensible wedges drive the timing. First, the Physical AI Bridge — a native ROS 2 + MCP + VLA orchestration layer for humanoid robots, cobots, and lab automation that no hyperscaler, no open-source framework, and no orchestration startup ships as of April 2026. Second, EU AI Act compliance — multi-agent orchestration in pharma, manufacturing, and financial services is classified high-risk under the Act, with enforcement beginning August 2026. Being the first orchestration fabric to achieve CE marking for high-risk AI adds $8–15M to each large enterprise implementation and compounds with every month of certification lead time.

This role is the most critical single hire gsiso will make. The Founding Engineer for Distributed Systems builds the Agent Mesh OS: the Ray-based scheduler, the LiteLLM multi-model routing layer, the agent lifecycle service, and the tenant isolation architecture that everything else runs on. You will be employee #1 (excluding the founder), you will make architectural decisions that cannot be undone, and you will be the person other engineers turn to when the distributed system misbehaves at 2am before a design partner demo. If you want to build infrastructure that governs physical robots, routes hundreds of concurrent agent threads, and satisfies EU AI Act audit requirements in production — this is the role.

What you'll build in your first 90 days

- Stand up the Ray 2.x cluster on AWS EKS (us-east-1 + eu-west-1), define the agent scheduling primitives, and ship the first passing CI pipeline on GitHub Actions — target: Sprint 1 complete by July 11
- Integrate LiteLLM as the multi-model routing gateway across OpenAI, Anthropic, Gemini, and a local vLLM inference endpoint; write the routing policy interface that lets OPA control which model an agent is allowed to call
- Design and ship the agent lifecycle service: agent spawn, heartbeat, graceful shutdown, and forced kill — with every lifecycle event producing an ed25519-signed audit receipt in the trust

ledger

• Build the tenant isolation layer on Postgres 16 + pgvector 0.7: per-tenant agent namespace, row-level security policies, and the semantic memory query interface that agent instances call at runtime

• Define the agent message contract schema (typed, versioned, and validated against a JSON Schema registry) so that the robotics engineer's MCP servers and the security engineer's DID layer have a stable interface to build against

• Establish the OpenTelemetry instrumentation baseline across all services: traces, metrics, and structured logs flowing to ClickHouse before the first design partner demo

• Prepare the Phase 1 integration test harness so the team can run a simulated 100-agent load test against the scheduler before Sprint 3's Claude→ROS 2 demo

What you'll own long-term

The Agent Mesh OS scheduler, multi-model routing layer, agent lifecycle service, horizontal scaling architecture, tenant isolation model, inter-agent message contract system, and the distributed systems reliability posture of the platform. As the team grows, you will hire and own the distributed systems sub-team. You will be the technical authority on every architectural decision that touches agent scheduling, memory, or message routing.

Required experience

• 5+ years building distributed systems in Python (async/asyncio) or Go — not just using distributed systems, but having designed them: schemas, failure modes, recovery paths

• Production experience with Ray, Celery, or a comparable distributed task scheduler at scale (>10K concurrent tasks, multi-node)

• Kubernetes (EKS or GKE) in production: you have written Terraform modules, debugged pod evictions, and reasoned about cluster autoscaling trade-offs under load

• Postgres in production at multi-tenant scale: row-level security, query plan analysis, index design — not just schema migration

• Redis 7 as a hot-state cache or work queue: you have reasoned about eviction policies, persistence modes, and the failure envelope when Redis is unavailable

• API surface design with FastAPI or equivalent: you have shipped a versioned REST/HTTP API that external systems depend on, not just internal tooling

• Familiarity with LLM API calling patterns (streaming, token budgets, retry logic) at the level of someone who has built a production wrapper around at least one model provider

Bonus

- Prior work on agent orchestration frameworks (LangGraph, CrewAI, OpenAI Agents SDK) at the infrastructure level — not just usage, but extension or internal tooling built on top
- Experience with pgvector or another vector search system in a production multi-tenant environment
- Familiarity with W3C DID specifications, Verifiable Credentials, or similar identity primitives — useful context for the trust ledger you will interface with
- Shipped a system under SOC2 Type II or equivalent audit requirements — knowing what "audit evidence" means from the inside
- Prior startup experience (Series A or earlier), specifically the experience of making architectural decisions without the luxury of reversibility

---

Compensation

Base salary: $180,000–$220,000 depending on experience Equity: 1.0–1.5% of gsiso.ai, grant at offer, 4-year vesting schedule with 1-year cliff Benefits: Full medical/dental/vision + 401(k) + home office equipment budget; details confirmed at offer Vesting: 4 years, 1-year cliff. Standard early-exercise option available.

How to apply

Email gaurav@gsiso.ai with:

- (a) your résumé
- (b) a link to a distributed system you shipped that you're proud of — ideally one that is running in production
- (c) one paragraph on why this problem matters to you

No cover letter.

What we don't care about

- We don't care if you didn't go to a top CS program or any CS program. We care about the systems you have shipped and can explain at the design level.
- We don't care about LeetCode performance. We do not administer algorithm puzzles. The technical interview is a design session about distributed systems and failure modes.
- We don't care if your most recent role was at a startup no one has heard of. If you built the infrastructure, that is the credential.

Founding Engineer — Robotics / ROS 2

Location: Bay Area, CA — Hybrid (3 days on-site) Employment type: Full-time Reports to: Gaurav Sisodia, Founder Hire sequence: #2 — Month 0 Compensation: $180,000–$220,000 base + 1.0–1.5% equity

Why now — the company in 60 seconds

gsiso.ai is building the agentic intelligence fabric for the physical + digital economy: a platform that orchestrates AI agents across robotics, pharma, manufacturing, and capital markets with cryptographic identity and EU AI Act-compliant governance built into every layer. The largest whitespace in the entire AI stack right now is the software bridge between LLM agents and physical robot systems. No orchestration platform — LangGraph, CrewAI, Azure AI Foundry, AWS Bedrock AgentCore — ships native ROS 2 integration, robot fleet telemetry ingest, or VLA policy update pipelines as of April 2026. The global robotics market hit $38B in 2026 at 34% YoY growth; VLA model adoption tripled in 2025–2026 and now backs 40% of new robot deployments. The bottleneck is not hardware and is not model capability: it is the orchestration fabric that manages, audits, and governs these physical systems.

This role owns the Physical AI Bridge — the most defensible pillar of the gsiso platform. You will author the ROS 2 MCP server adapters for each robot class we support (UR10e cobot, OT-2 lab arm, MiR250 AMR, humanoid), build the WebRTC telemetry ingest pipeline, implement the Isaac Sim gate that validates every VLA policy update before physical deployment, and ship the safe-stop proof primitive that satisfies EU AI Act human oversight requirements. This is a hands-on engineering role in a lab with physical hardware from month 3 onward. If you have spent your career in robotics software and are frustrated that the AI orchestration world has ignored physical systems, this is the role to fix that.

What you'll build in your first 90 days

- Stand up ROS 2 Jazzy in a Docker container with a stable CI test harness; author the first MCP server skeleton against the ROS 2 action client interface — deliverable for Sprint 1
- Ship the UR10e MCP server: translate LLM agent intent (via MCP tool calls) to ROS 2 action goals for the UR10e cobot; validate the full loop in NVIDIA Isaac Sim before any physical hardware is involved — deliverable for Sprint 3
- Build the OT-2 lab arm integration: OT-2 MCP server using the Opentrons Python API over the ROS 2 bridge; execute a real 3-step pipetting protocol driven by a GPT-4o agent — deliverable for Sprint 4

- Implement the safe-stop proof primitive: a triple-signed token (agent DID + operator key + hardware attestation) generated at halt-command time and validated by the robot's safety controller; p99 safe-stop latency target is ≤ 120ms — deliverable for Sprint 5
- Implement the WebRTC telemetry ingest pipeline for real-time robot state streaming (joint angles, end-effector pose, tool state, error flags) from robot fleet to the gsiso control plane — deliverable shared with Sprint 4 and Sprint 6
- Build the Isaac Sim gate pipeline: every VLA policy update must execute a full sim run against a representative task set and pass a configurable success threshold before the update is deployed to physical hardware
- Begin the humanoid MCP server (1X NEO or Figure 02 / GR00T N1 dispatch path) in Sprint 7, with a full Isaac Sim demo of an LLM-directed humanoid action sequence complete before month 6

What you'll own long-term

The Physical AI Bridge in its entirety: all per-robot-class MCP server adapters (humanoid, cobot, AMR, lab arm, drone), the sim-to-real deployment pipeline, WebRTC telemetry ingest, safe-stop proof primitive, VLA policy management interface, and the physical hardware lab in our Bay Area office. As the team grows to include a second robotics engineer at M5, you will define the integration architecture they operate within and review their robot-class additions. You are accountable for physical system safety — every command that reaches a robot traces through your code.

Required experience

- 5+ years writing production ROS 2 (or ROS 1 with ROS 2 migration experience): you have shipped software that ran on physical robots in environments where failure had real consequences, not just demos
- C++ and Python fluency: ROS 2 node authoring requires both; the MCP server layer is Python; the safety-critical primitives may require C++
- Hands-on integration experience with at least one commercial robot platform — UR cobot (UR5/UR10/UR10e), Fanuc CR-series, KUKA iiwa, OT-2, or equivalent: you have written drivers, handled error states, and debugged physical hardware failures
- ROS 2 action server/client architecture: you understand preemption, feedback, and result handling under network partition conditions
- Simulation experience with NVIDIA Isaac Sim, Gazebo, or MuJoCo: you have used sim as a validation gate, not just a demonstration environment
- Familiarity with robot safety standards (ISO 10218, ISO/TS 15066 for cobots, or equivalent): you understand what safe-stop requirements mean at the hardware level and can translate them to software guarantees

Bonus

- Prior work building or integrating VLA models (π0, GR00T N1, Gemini Robotics, or academic equivalents) into a physical deployment pipeline
- Experience with the Opentrons OT-2 or similar liquid handling robot platform
- WebRTC implementation experience for real-time telemetry or video streaming from edge devices
- Familiarity with MCP (Model Context Protocol) or any agent tool-calling interface
- Prior startup experience, specifically having owned a system end-to-end in a team of fewer than 10 people
- EU AI Act or functional safety (IEC 61508) familiarity — knowing what "human oversight" means at the system level

---

Compensation

Base salary: $180,000–$220,000 depending on experience Equity: 1.0–1.5% of gsiso.ai, grant at offer, 4-year vesting schedule with 1-year cliff Benefits: Full medical/dental/vision + 401(k) + home office equipment budget; details confirmed at offer Vesting: 4 years, 1-year cliff. Standard early-exercise option available.

How to apply

Email gaurav@gsiso.ai with:

- (a) your résumé
- (b) a link to a system you shipped that you're proud of — ideally a physical robot integration or a production ROS 2 deployment
- (c) one paragraph on why this problem matters to you

No cover letter.

What we don't care about

- We don't care if your most impressive work was at a company that failed. Physical AI infrastructure is hard to ship; the failures are instructive, not disqualifying.
- We don't care about LeetCode performance. Our technical interview is a systems design session and a ROS 2 architecture walkthrough — we will talk through a real integration problem.
- We don't care if you have never worked with LLM APIs. You will learn the MCP protocol in days; we care that you can guarantee the safety of a ROS 2 action under adversarial inputs.

Founding Engineer — Security / Cryptography

Location: Bay Area, CA — Hybrid (3 days on-site) Employment type: Full-time Reports to: Gaurav Sisodia, Founder Hire sequence: #3 — Month 1 Compensation: $190,000–$230,000 base + 0.75–1.25% equity

Why now — the company in 60 seconds

gsiso.ai is building the trust and governance layer for the agentic economy. Our platform orchestrates AI agents across physical and digital systems — robots, financial infrastructure, pharma pipelines — and every agent action must be cryptographically attributable, tamper-evident, and provably compliant with EU AI Act Annex VIII technical documentation requirements. The Trust & Governance pillar is not a compliance checkbox on top of the product: it is the product's central value proposition. Enterprises in regulated industries pay $8–15M per large implementation to satisfy EU AI Act high-risk AI requirements; the ability to hand a Notified Body a Merkle-chained, ed25519-signed audit receipt log for every agent action in a production system is a direct procurement shortcut for European buyers — the same regulatory moat that made Veeva the mandatory cloud for pharma.

This role owns the cryptographic identity and trust ledger of the entire platform. You will build the agent DID minting service (W3C DID:web method), the ed25519 signing pipeline that produces tamper-evident audit receipts for every agent action, the Merkle chain that binds receipts into a verifiable log, the HashiCorp Vault integration for key lifecycle management and HSM support, and the kill switch propagation system that bypasses software to halt a physical robot or agent process. Anthropic's published research documented that every frontier model tested attempted to circumvent shutdown commands — the kill switch you build is not a theoretical safety primitive; it is a production-grade, mission-critical component. If you care about getting cryptographic identity and audit trail systems right in a world where the stakes are measured in robot safety and regulatory approval, this is the right place to do it.

What you'll build in your first 90 days

- Design and ship the agent DID minting service: generate a W3C DID:web identity for each agent instance at spawn time; store the DID document in Postgres with full version history; expose a resolution endpoint compliant with the DID:web spec
- Build the ed25519 signing pipeline: every significant agent action (robot command issued, policy evaluated, workflow modified, kill switch invoked) produces a signed audit receipt; receipts are serialized to a defined schema and written to the trust ledger in Postgres within 50ms of the event — Sprint 2 deliverable

- Implement Merkle chaining of audit receipts: each receipt includes the hash of the previous receipt in its signed payload; compute and commit the chain root to S3 on a configurable interval; root commitment is itself signed
- Integrate HashiCorp Vault for key lifecycle management: agent signing keys generated and stored in Vault, with rotation policy and HSM-backed storage path for production; write the Vault integration so that the signing service never holds private key material in application memory
- Implement Sigstore/Cosign signing for OCI agent images: every agent container image is signed at build time; the gsiso scheduler verifies the Cosign signature before instantiating any agent — eliminates unsigned agent execution
- Build the kill switch propagation system: a signed kill-switch command issued from the console or API must propagate to all running instances of an agent (and, via the robotics bridge, to any connected physical system) within a target latency; propagation is recorded in the trust ledger with a revocation receipt
- Author the initial trust ledger query API so that the compliance lead (joining at M3) can extract Annex VIII technical documentation evidence from the ledger without writing SQL

What you'll own long-term

The Trust & Governance plane in its entirety: agent DID minting service, ed25519 signing pipeline, Merkle-chained trust ledger, Vault key lifecycle integration, Sigstore/Cosign agent image signing, kill switch propagation, and the cryptographic policy enforcement layer that underpins the EU AI Act evidence pipeline. As the team grows, you are the on-call authority for any question that begins "how do we prove that this agent did X at time T." You will work directly with the Compliance Lead (hire #6) to define what the trust ledger must produce to satisfy Notified Body requirements.

Required experience

- 5+ years in security engineering, cryptographic systems, or PKI infrastructure — not just consuming crypto libraries, but designing the systems that produce and validate cryptographic proofs
- Production experience with ed25519, ECDSA, or comparable asymmetric signing schemes: you have shipped a signing pipeline, managed key material, and reasoned about signing failure modes at scale
- HashiCorp Vault in production: secrets management, dynamic credentials, key/value store with access policies — you have built an integration, not just read the docs
- Familiarity with W3C DID specifications (DID:web, DID:key, or similar): you understand the resolution protocol, document format, and key rotation mechanics
- PKI system design: certificate issuance, revocation (CRL or OCSP), chain validation — you understand the trust anchors and the failure modes when a link in the chain is compromised

- Go or Rust experience at the production level: the cryptographic core will be implemented in one of these; Python is not acceptable for the signing service due to GIL and memory safety considerations

Bonus

- Experience with Sigstore, Cosign, or supply chain security tooling (SLSA, in-toto)
- Familiarity with Merkle trees or certificate transparency logs (RFC 6962): you have either implemented one or audited an implementation
- Prior work in a regulated environment (SOC2 Type II, FedRAMP, ISO 27001, or HIPAA) where audit trails were a hard compliance requirement, not a nice-to-have
- Experience with hardware security modules (HSMs): Thales Luna, AWS CloudHSM, or equivalent — gsiso will require HSM-backed key storage for enterprise customers
- Familiarity with EU AI Act Annex VIII technical documentation requirements or NIST AI RMF 2.0 — useful context for designing what the trust ledger must capture
- Prior experience at a security startup (Series A or earlier) or a cryptographic infrastructure role at a company where the identity layer was the product

Compensation

Base salary: $190,000–$230,000 depending on experience Equity: 0.75–1.25% of gsiso.ai, grant at offer, 4-year vesting schedule with 1-year cliff Benefits: Full medical/dental/vision + 401(k) + home office equipment budget; details confirmed at offer Vesting: 4 years, 1-year cliff. Standard early-exercise option available.

How to apply

Email gaurav@gsiso.ai with:

- (a) your résumé
- (b) a link to a system you shipped that you're proud of — ideally one where the cryptographic or identity layer is the part you are most satisfied with
- (c) one paragraph on why this problem matters to you

No cover letter.

What we don't care about

- We don't care if you have a background in financial security, infrastructure security, or IoT security — the primitives transfer and we care more about your depth in cryptographic system design than your industry of origin.
- We don't care about LeetCode performance. Our technical interview is a design session: we will walk through a signing pipeline failure mode and ask how you would detect and recover from it.
- We don't care if your opinion on DID:web versus DID:key differs from ours — we made a choice for defensible reasons and we will explain it; if your argument is better, we want to hear it.

Design Engineer — Console + Policy Studio UI

Location: Bay Area, CA — Hybrid (3 days on-site) Employment type: Full-time Reports to: Gaurav Sisodia, Founder Hire sequence: #4 — Month 1 Compensation: $150,000–$185,000 base + 0.5–1.0% equity

Why now — the company in 60 seconds

gsiso.ai is building the operating layer for the agentic economy: a platform that orchestrates AI agents across robot fleets, pharma pipelines, and capital markets with cryptographic identity and EU AI Act-compliant governance at every layer. The control plane console is not a dashboard bolted onto a backend — it is the primary product surface that enterprise buyers evaluate, compliance officers audit against, and operators use to watch a robot pick up a vial while a signed receipt appears in real time. The Policy Studio — a UI layer for authoring, testing, and deploying Rego-based access policies — is a commercial moat: it makes our OPA-backed policy engine accessible to compliance engineers who cannot write Rego, and it is the feature most cited in enterprise sales discussions as a differentiation point against the open-source alternatives.

This role owns the entire frontend surface of gsiso. You will build the console from a clean codebase using Next.js 15, React 19, Tailwind CSS 4, and shadcn/ui; integrate Clerk for OIDC/SAML enterprise auth; wire up tRPC for end-to-end type-safe APIs to the backend; and ship the Policy Studio as a first-class product by Phase 2. This is a role for someone who thinks about design and engineering together — who is frustrated by enterprise software that is technically complete and visually incoherent, and who wants to build something that compliance engineers and robotics operators actually prefer to use. The console is the face of gsiso to every enterprise customer we have. The quality of what you build directly determines whether we close design partners.

What you'll build in your first 90 days

- Stand up the Next.js 15 App Router codebase with Tailwind CSS 4, shadcn/ui, Clerk auth (OIDC federation to enterprise IdP), Stripe billing metering integration, and a tRPC client wired to the backend API — clean, typed, tested, documented, and production-deployable by Sprint 1
- Ship the agent activity dashboard: real-time view of running agents (DID, status, current task, model being called, policy state), with a live trust receipt feed showing each signed audit receipt as it is written to the ledger — this is the demo screen for every design partner meeting
- Build the fleet telemetry view: WebRTC stream ingest from robot fleet displayed in a browser-based UI, showing joint state, end-effector pose, error flags, and the last 10 signed commands for each connected robot — Sprint 4 deliverable used in the OT-2 demo

- Ship the kill switch UI: a single, high-confidence interaction surface that issues a hardware-level halt command to a selected agent or robot; the UI must display the propagation confirmation and the resulting revocation receipt before allowing dismissal
- Begin the Policy Studio alpha: a Rego policy authoring interface with syntax highlighting, inline evaluation feedback against a test input set, a diff view for policy changes, and a deployment pipeline that commits the policy to the OPA policy store — alpha shipped before the Phase 2 design partner onboarding
- Implement the Policy Studio compliance view: an audit export surface that packages selected trust receipts, agent DID resolution records, and policy evaluation logs into a downloadable evidence bundle structured for EU AI Act Annex VIII documentation — a feature no competing console product offers

What you'll own long-term

The full console product surface: agent dashboard, fleet telemetry views, Policy Studio, compliance evidence export, onboarding flows, billing and metering surfaces, and the design system that governs visual consistency across all of them. You own the design system definition, the component library built on shadcn/ui primitives, and the engineering quality of every user-facing interaction. As the company grows, you will hire the frontend team and define the standards they work within. The Policy Studio is your flagship product and the one that most directly differentiates gsiso from every competing platform.

Required experience

- 5+ years building production enterprise UI in React — not consumer apps: dashboards, admin consoles, data-dense compliance interfaces where accessibility, keyboard navigation, and visual hierarchy under load matter
- Next.js (App Router, not just Pages Router): you understand server components, server actions, caching semantics, and the deployment model; you have shipped a Next.js app to production and reasoned about its performance characteristics
- TypeScript: you write typed frontend code as a default, not a project requirement; you understand how tRPC or similar end-to-end type safety changes the development loop
- Tailwind CSS and a component library (shadcn/ui, Radix UI, or equivalent): you have built a design system on these foundations, not just used preset components
- Real-time UI engineering: WebSockets, SSE, or WebRTC data channel consumption in a production context — displaying live data without degrading page performance
- Enterprise auth integration: Clerk, Auth0, or direct OIDC/SAML — you have wired up enterprise IdP federation and understand RBAC at the UI layer

Bonus

- Prior work on a policy authoring or code editor UI (Monaco, CodeMirror, or similar) — the Policy Studio's Rego authoring interface requires embedded code editing
- Experience building compliance or audit dashboards in regulated industries (finance, healthcare, or industrial) — familiarity with what "audit evidence" means to the person who will use the export feature
- Design skills: you can produce a mid-fidelity Figma mockup and move it to code without a handoff process; you own the visual quality of what you ship
- Familiarity with Stripe Billing metering and usage-based pricing UI patterns
- Prior experience at a B2B SaaS company where the UI was a key differentiator against open-source alternatives (Grafana vs. open-source dashboarding; Retool vs. internal tool builders; HashiCorp Vault UI vs. CLI analogues)

Compensation

Base salary: $150,000–$185,000 depending on experience Equity: 0.5–1.0% of gsiso.ai, grant at offer, 4-year vesting schedule with 1-year cliff Benefits: Full medical/dental/vision + 401(k) + home office equipment budget; details confirmed at offer Vesting: 4 years, 1-year cliff. Standard early-exercise option available.

How to apply

Email gaurav@gsiso.ai with:

- (a) your résumé
- (b) a link to a system you shipped that you're proud of — a link to a production UI you built or a portfolio that shows the quality of your work; screenshots are fine if the product is behind auth
- (c) one paragraph on why this problem matters to you

No cover letter.

What we don't care about

- We don't care if you have a design degree or an engineering degree or neither. We care about the quality of what you have shipped and your ability to defend the decisions you made.
- We don't care about LeetCode performance. Our interview is a design and code review: we will look at code you have written and UI you have built and ask you to walk us through your decisions.

- We don't care if you have never worked in robotics or compliance — the context transfers quickly and the UI engineering problem is the same whether the data on screen is a robot joint angle or a financial instrument. We care about enterprise UI depth, not domain familiarity.

SECTION 05

Design Partner Outreach

Three primary targets: Recursion (pharma wet-lab automation), Siemens (industrial AI orchestration + EU AI Act friction), Two Sigma (data-as-code research workflows). Seven backup targets in the tracker. Outreach goes out the week of May 18.

Design-Partner Outreach Tracker — gsiso.ai Phase 0

Owner: Gaurav Sisodia Last updated: May 2026 Target: 3 signed LOIs by end of Phase 0 (June 2026) LOI contract shape: 12-month fee-waived pilot (first anchor per vertical) + defined success metrics + joint press release rights at GA + customer reference rights at Series A

Primary Targets (emails drafted)

Target	Contact Name (best guess)	LinkedIn Search Query	Vertical	Sent Date	Status	Last Touch	Next Action	Notes
Recursion Pharmaceuticals	*Search required*	`"Recursion Pharmaceuticals" "Head of Platform Engineering" OR "VP Platform"`	Pharma / Biotech	—	Draft ready	—	Send Mon May 12; follow up May 19 if no reply	Success metric: 10 compound prep protocols/week with full signed audit trail, zero manual re-runs, in 90 days. First pharma anchor → fee-waived pilot. Email: `01-recursion-pharmaceuticals.md`

Target	Contact Name (best guess)	LinkedIn Search Query	Vertical	Sent Date	Status	Last Touch	Next Action	Notes
Siemens Digital Industries	*Search required*	`"Siemens Digital Industries" "Head of Industrial AI" OR "VP Industrial AI" OR "Director Industrial AI"`	Manufacturing / Industrial	—	Draft ready	—	Send Tue May 13; follow up May 20 if no reply	Success metric: UR10e cobot fleet under gsiso governance at one Siemens plant; VLA policy update via sim gate without downtime. Cedrik Neike (CEO DI) is too senior — go one level below. Email: `02-siemens-digital-industries.md`

Target	Contact Name (best guess)	LinkedIn Search Query	Vertical	Sent Date	Status	Last Touch	Next Action	Notes
Two Sigma	*Search required*	`"Two Sigma" "Head of Systems" OR "Head of Infrastructure" OR "VP Engineering" site:linkedin.com`	Capital Markets	—	Draft ready	—	Send Wed May 14; follow up May 21 if no reply	Contact name gap — Two Sigma does not publish org chart. Consider warm intro via ICML/NeurIPS network before cold send. Success metric: internal agent workflow across 2 model providers, MiFID II audit trails, under gsiso governance. Email: `03-two-sigma.md`

Additional Recommended Targets (5–7)

Target	Contact Name (best guess)	LinkedIn Search Query	Vertical	Sent Date	Status	Last Touch	Next Action	Notes
Insitro	*Search required*	`"Insitro" "CTO" OR "Head of Engineering" OR "Head of Infrastructure"`	Pharma / Biotech	—	Not started	—	Queue for M1 if Recursion LOI not closed	ML-first drug discovery startup; smaller decision-making unit than Big Pharma; familiar with ROS 2-adjacent lab automation. Delivery plan §7 identifies as secondary pharma target. Success metric: Insitro lab automation pipeline runs through gsiso Physical AI Bridge; audit receipts satisfy internal compliance review.

Target	Contact Name (best guess)	LinkedIn Search Query	Vertical	Sent Date	Status	Last Touch	Next Action	Notes
Fanuc America	*Search required*	`"Fanuc America" "Chief Robotics Officer" OR "VP Software" OR "Head of Software Platforms"`	Manufacturing / Industrial	—	Not started	—	Queue for M1	Largest cobot installed base in North America (CR-series). No existing AI orchestration layer — greenfield opportunity. Success metric: five Fanuc CR-series cobots receive LLM-directed task scheduling via gsiso MCP bridge with measurable throughput improvement in 30 days. Delivery plan §7.

Target	Contact Name (best guess)	LinkedIn Search Query	Vertical	Sent Date	Status	Last Touch	Next Action	Notes
Citadel Securities	*Search required*	`"Citadel Securities" "Head of Technology Infrastructure" OR "VP Technology" OR "Head of Platform"`	Capital Markets	—	Not started	—	Queue for M1 after Two Sigma outreach is processed	Multi-cloud, compliance-obsessed; high ACV potential. Success metric: research synthesis agent swarm producing MiFID II-compliant audit trails; agent actions traceable to individual DID in < 5 seconds. Delivery plan §7.
Boston Dynamics (Hyundai)	*Search required*	`"Boston Dynamics" "VP Software Platforms" OR "Head of Software" OR "Director Software Engineering"`	Manufacturing / Industrial (Humanoid)	—	Not started	—	Queue for M2 after Phase 1 humanoid demo is complete	Atlas pilot at Hyundai Georgia plant is live; software orchestration is the gap. Approach after Sprint 7 humanoid demo provides credible demo artifact. Delivery plan §7.

Target	Contact Name (best guess)	LinkedIn Search Query	Vertical	Sent Date	Status	Last Touch	Next Action	Notes
Novartis Digital & Tech	*Search required*	`"Novartis" "Head of AI Lab" OR "Head of Digital Innovation" Basel`	Pharma / Biotech (Enterprise)	—	Not started	—	Queue for M2 after EU AI Act compliance narrative is sharpened	EU HQ validates EU AI Act story; large compliance budget. Scale validates enterprise readiness. Pilot: Lead-to-IND agent pack with physical lab arm; at least one EU AI Act Annex III requirement demonstrably satisfied. Delivery plan §7.
Vericel / AgBiome / Recursion spinout or peer	*Research required*	`TechBio "Head of Platform" OR "Head of Engineering" site:linkedin.com`	Pharma / Biotech (Backup)	—	Not started	—	Queue for M1 if Recursion and Insitro both decline	Backup pharma design partner if primary two do not convert. Target: any AI-native biotech with a robotic lab platform and ≤2 month procurement cycle. Self-sourced via robotics PhD network outreach.

Target	Contact Name (best guess)	LinkedIn Search Query	Vertical	Sent Date	Status	Last Touch	Next Action	Notes
D.E. Shaw Research	*Search required*	`"D.E. Shaw Research" "Head of Infrastructure" OR "Head of Engineering" OR "VP"`	Capital Markets / Computational Science	—	Not started	—	Queue for M2	Computationally sophisticated; known for custom infrastructure; smaller decision-making unit than Two Sigma or Citadel. Relevant if Two Sigma pilot is slow to close.

Cadence

- Week 1 (May 12–16): Send Recursion (Mon), Siemens (Tue), Two Sigma (Wed). No more than 2 new cold contacts per week to preserve reply bandwidth.
- Week 2 (May 19–23): Follow up on any non-replies from Week 1. Queue Insitro and Fanuc for send.
- Week 3 (May 26–30): Follow up Week 2. Queue Citadel for send.
- M1 (June): Activate remaining targets based on response rate from primary three.
- M2 (July): Boston Dynamics and Novartis activated after Phase 1 sprint 3–4 demo artifacts are available — a credible demo is the single most effective cold outreach asset.

Status definitions

`Draft ready` | `Sent` | `Opened (no reply)` | `Replied` | `Call scheduled` | `Call complete` | `LOI in negotiation` | `LOI signed` | `Declined` | `Ghosted — follow up`

Design-Partner Outreach — Recursion Pharmaceuticals

Email

Subject: Agent orchestration + audit trails for Recursion's robotic lab

From: Gaurav Sisodia \<gaurav@gsiso.ai\> To: Head of Platform Engineering, Recursion Pharmaceuticals

I'm Gaurav Sisodia, founder of gsiso.ai — I'm writing because Recursion is the only pharma company I know of where the platform engineering team owns both the ML infrastructure and the robotic lab automation layer simultaneously, which is exactly the integration surface we are building for.

gsiso.ai is the orchestration fabric that connects LLM agents to physical lab systems — specifically, ROS 2-based lab arms and liquid handling robots — with cryptographic audit trails on every agent action and EU AI Act-compliant governance built in. Recursion's platform engineering blog has documented the challenge of closing the loop between autonomous lab operations and traceable, auditable records; your computer vision work for autonomous error correction in the robotic operating system (December 2024) is precisely the kind of multi-step physical agent workflow our Physical AI Bridge is designed to govern, audit, and replay without manual intervention.

What we are proposing is a 12-month design-partner pilot, fee-waived for the first pharma anchor. The success metric is specific: an LLM agent + OT-2 lab arm completes 10 compound preparation protocols per week with a full, cryptographically signed audit trail — zero manual re-runs required — within 90 days of integration. We write the integration spec jointly, you define the protocol; we provide the orchestration, audit receipts, and the kill switch that satisfies any internal compliance review.

Would you have 30 minutes the week of May 18 to walk through what this looks like in your lab environment? I can make Tuesday May 19 at 10am PT, Wednesday May 20 at 2pm PT, or Thursday May 21 at 11am PT work.

Gaurav Sisodia Founder, gsiso.ai https://gsiso.ai | Delivery plan: https://gsiso.ai/delivery-plan.pdf

P.S. The State of Robotics 2026 report found that VLA model adoption tripled in 2025–2026 and now backs 40% of new robot deployments — but the audit and governance infrastructure for those deployments is the largest unaddressed gap in the stack. That pattern maps closely to where Recursion's lab infrastructure sits today.

Notes

Cited source:

- Recursion blog post "Improving Wet Lab Automation with Computer Vision" (December 4, 2024): https://www.recursion.com/news/improving-wet-lab-automation-with-computer-vision-a-hack-week-breakthrough — confirms Recursion is integrating computer vision into its robotic OS for autonomous error detection and correction, directly relevant to the gsiso Physical AI Bridge value proposition.
- Recursion OS platform page: https://www.recursion.com/platform — confirms robotics + ML infrastructure ownership in a single platform layer.

Research before they reply:

- Find the current Head of Platform Engineering at Recursion on LinkedIn (search: "Recursion Pharmaceuticals" + "Platform Engineering" + "Head" or "VP"). The role may have changed; verify the name before sending.
- Review Recursion's most recent engineering blog posts and any public talks at NeurIPS/ICML about their infrastructure — tailor the technical depth of the follow-up accordingly.
- Check whether Recursion has announced any EU lab expansion (they have EMEA operations) — EU AI Act angle may be more or less relevant depending on their lab footprint.
- Identify Recursion's internal compliance workflow for GxP and FDA 21 CFR Part 11 — this is the regulatory framing they care about, which maps to our audit receipt architecture.

Success looks like:

- Response rate target: 20–30% for this vertical (pharma platform engineering is a niche audience; a warm referral would double this).
- Next-step CTA: 30-minute video call, week of May 18. If they respond positively but push timing, offer a 15-minute async Loom walkthrough of the sprint plan as a low-commitment alternative.
- The design partner LOI must be signed by end of Phase 0 (June 2026). A first response in May gives adequate runway for a 2–3 week LOI negotiation cycle.

Design-Partner Outreach — Siemens Digital Industries

Email

Subject: Vendor-neutral agent orchestration for Siemens cobot fleets

From: Gaurav Sisodia \<gaurav@gsiso.ai\> To: Head of Industrial AI, Siemens Digital Industries

I'm Gaurav Sisodia, founder of gsiso.ai — I'm reaching out because Siemens Digital Industries has built exactly the kind of open, multi-provider industrial AI orchestration architecture (WinCC OA MCP Server, multi-LLM integration across OpenAI, AWS Bedrock, Claude, and Gemini) that surfaces the governance gap we are designed to close.

gsiso.ai is the vendor-neutral orchestration fabric that adds cryptographic identity, tamper-evident audit trails, and EU AI Act-compliant governance to the layer between LLM agents and physical automation systems. Your industrial AI orchestration layer documentation describes the requirement precisely: AI must be "safe, deterministic, and auditable" — with "full traceability" and validated model updates deployed via virtual commissioning before they reach a PLC. That is the architecture we ship as a managed platform, with a trust ledger that produces Annex VIII technical documentation evidence automatically. We are aware that Siemens has raised concerns about the EU AI Act's treatment of industrial AI applications as equivalent to consumer applications; our platform is designed to make compliance evidence generation so low-friction that it removes the regulatory overhead objection rather than adding to it.

We are proposing a 12-month design-partner pilot, fee-waived for the first manufacturing anchor. The success metric: a UR10e cobot fleet at one Siemens facility managed through the gsiso Physical AI Bridge, with VLA policy updates executed via our Isaac Sim gate without production downtime — measurable in 30 days of operation. We write the integration spec jointly against your existing ROS 2 and OPC UA infrastructure.

I can make Tuesday May 19 at 9am PT / 6pm CET, Wednesday May 20 at 3pm PT / midnight CET, or Thursday May 21 at 8am PT / 5pm CET work for a 30-minute call the week of May 18.

Gaurav Sisodia Founder, gsiso.ai https://gsiso.ai | Delivery plan: https://gsiso.ai/delivery-plan.pdf

P.S. The Siemens Industrial AI Orchestration Layer architecture document describes the challenge of resolving conflicts between multiple AI systems before anything reaches a PLC — that conflict-resolution problem at scale, across an agent mesh rather than a single copilot, is where the gsiso policy engine applies directly.

Notes

Cited sources:

- Siemens Industrial AI Orchestration Layer page: https://www.siemens.com/en-us/company/insights/industrial-operations-x/architecture-hub/industrial-ai-orchestration-layer/ — directly references multi-LLM provider integration (OpenAI, AWS Bedrock, Claude, Gemini) via WinCC OA MCP Server, and the requirement for "safe, deterministic, and auditable" AI with "full traceability."
- Bloomberg / Techzine reporting (April 20–21, 2026) on Siemens CEO Roland Busch's statement that the EU AI Act and Data Act "miss the mark" by treating industrial AI like consumer applications: https://www.techzine.eu/news/infrastructure/140602/siemens-warns-eu-regulations-are-slowing-down-ai-investments/ — the P.S. and body framing address this directly.

Research before they reply:

- Find the current Head of Industrial AI at Siemens Digital Industries on LinkedIn (search: "Siemens Digital Industries" + "Head of Industrial AI" or "VP Industrial AI"). Cedrik Neike is CEO of Siemens Digital Industries but is not the right contact for a pilot conversation — the relevant contact is one or two levels below in the industrial AI product or R&D organization.
- Review the Eigen Engineering Agent announcement from Hannover Messe — this is Siemens' most recent move toward autonomous execution in industrial AI and is the direct adjacent capability to what gsiso governs. Reference it in the follow-up if they respond.
- Clarify whether Siemens Digital Industries has an internal ROS 2 deployment (they have public documentation referencing ROS 2 integration in their factory automation stack) — tailor the technical integration depth in follow-up accordingly.
- Understand Siemens' Xcelerator partner ecosystem model — gsiso may fit as a certified Xcelerator partner rather than (or in addition to) a design partner, which is a more scalable distribution path. Raise this in the second meeting if the first goes well.

Success looks like:

- Response rate target: 15–25%. Siemens is a large organization; the cold email will likely be forwarded internally before it reaches the right person. That is acceptable.
- Next-step CTA: 30-minute call, week of May 18. The goal of the first call is not to close an LOI — it is to identify whether there is a specific plant, product team, or internal AI initiative with the right combination of ROS 2 infrastructure and EU AI Act compliance pressure to make a pilot tractable within 6 months.
- Siemens LOI timing is likely M1–M2 rather than M0 given their procurement cycle.

Design-Partner Outreach — Two Sigma

Email

Subject: Multi-provider agent governance without single-model lock-in

From: Gaurav Sisodia \<gaurav@gsiso.ai\> To: Head of Systems, Two Sigma

I'm Gaurav Sisodia, founder of gsiso.ai — I'm reaching out because Two Sigma's public writing on data infrastructure and ML systems signals an organization that builds everything with a systems-design discipline uncommon in the financial industry, and the agent governance problem we solve sits directly in that stack.

gsiso.ai is the vendor-neutral orchestration fabric for enterprise AI agents: cryptographic identity (ed25519 DIDs) for every agent, tamper-evident audit receipts on every action, and a policy engine that controls which agents can call which models — all without locking your infrastructure to a single provider. Two Sigma's published approach to data infrastructure — version-controlled pipelines, CI/CD for data transformations, and sophisticated internal orchestration systems for complex computational workflows — maps directly to the architecture our Agent Mesh OS extends into the agentic layer. The specific capability we bring is multi-provider agent governance: two LLM providers running concurrently under a single policy enforcement point, with every inter-agent action traceable to a specific DID in under five seconds.

We are proposing a 12-month design-partner pilot, fee-waived for the first capital markets anchor. The success metric is concrete: a Two Sigma internal agent workflow running across two model providers simultaneously under gsiso governance, with no single-provider lock-in and MiFID II-grade audit trails on every agent decision — production-verifiable within 60 days of integration start. We write the integration spec to fit your existing infrastructure; we do not require you to replace anything.

Would 30 minutes the week of May 18 work to walk through how this fits your stack? I can make Monday May 18 at 1pm ET, Wednesday May 20 at 10am ET, or Friday May 22 at 2pm ET work.

Gaurav Sisodia Founder, gsiso.ai https://gsiso.ai | Delivery plan: https://gsiso.ai/delivery-plan.pdf

P.S. Two Sigma's 2025 ICML coverage highlighted sequential anytime-valid inference and continuous monitoring under optional stopping — the same statistical rigor applied to agent behavior evaluation (outcome scoring across concurrent model providers) is a design challenge we are working through in our evaluation harness; I would be curious whether that maps to anything on your end.

Notes

Cited sources:

- Two Sigma "Treating Data as Code" article (November 2025): https://www.twosigma.com/articles/treating-data-as-code-at-two-sigma/ — references "sophisticated orchestration systems for managing complex computational workflows" and CI/CD-style pipeline governance, directly analogous to what gsiso extends into the agent layer.
- Two Sigma ICML 2025 recap (July 2025): https://www.twosigma.com/articles/icml-2025-key-ideas-on-llms-human-ai-alignment-and-more/ — confirms active ML research engagement; P.S. references sequential anytime-valid inference (SAVI) coverage from their attendees as a specific common-ground signal.
- Two Sigma open source page: https://www.twosigma.com/open-source/ — confirms vendor-neutral, open-source-friendly infrastructure philosophy consistent with gsiso's model-neutral positioning.

Research before they reply:

- Contact name gap: A specific named Head of Systems at Two Sigma is not publicly identifiable with confidence. LinkedIn search recommended: query "Two Sigma" + "Head of Systems" or "VP Engineering" or "Head of Infrastructure." Relevant titles include Chief Technology Officer, Head of Core Infrastructure, or Head of Platform Engineering. Two Sigma is known to be private about its organizational structure; a warm introduction via a mutual contact (e.g., someone from the quantitative finance or ML infrastructure community) is meaningfully more effective than a cold email here.
- Review Two Sigma's recent open-source contributions (twosigma.com/open-source) for any agent orchestration or workflow management projects — if they have internal tooling in this space, position gsiso as a governance layer on top rather than a replacement.
- Check whether Two Sigma has any public statements on MiFID II compliance tooling or agent-based trading system governance — this strengthens the audit trail angle in the follow-up.
- Two Sigma Ventures (the VC arm) is separate from Two Sigma the hedge fund; ensure outreach is directed to the operating company, not the investment arm.

Success looks like:

- Response rate target: 10–20%. Two Sigma is notoriously internal; cold outreach has lower baseline response rates than the pharma or manufacturing targets. A warm introduction via the ML research community (NeurIPS, ICML mutual contacts) could raise this to 30–40%.
- Next-step CTA: 30-minute call, week of May 18. If no response within 7 days, a single follow-up is appropriate; more than two contacts without a response suggests the outreach channel is wrong, not the message.
- Alternative path: Two Sigma Ventures has published publicly about multi-provider LLM orchestration and agent ETL pipelines (2024 predictions piece). If the operating company does not respond, Two Sigma Ventures may be a viable warm introduction path to the systems team.

END OF PACKET

gsiso.ai

The Agentic Intelligence Fabric

Owner: Gaurav Sisodia · gaurav@gsiso.ai
gsiso.ai · Internal — not for distribution